

Partially Linear Estimation with Application to Sparse Signal Recovery From Measurement Pairs

Tomer Michaeli, Daniel Sigalov and Yonina C. Eldar, *Senior Member, IEEE*

Abstract—We address the problem of estimating a random vector X from two sets of measurements Y and Z , such that the estimator is linear in Y . We show that the partially linear minimum mean squared error (PLMMSE) estimator does not require knowing the joint distribution of X and Y in full, but rather only its second-order moments. This renders it of potential interest in various applications. We further show that the PLMMSE method is minimax-optimal among all estimators that solely depend on the second-order statistics of X and Y . We demonstrate our approach in the context of recovering a signal, which is sparse in a unitary dictionary, from noisy observations of it and of a filtered version of it. We show that in this setting PLMMSE estimation has a clear computational advantage, while its performance is comparable to state-of-the-art algorithms. We apply our approach both in static and dynamic estimation applications. In the former category, we treat the problem of image enhancement from blurred/noisy image pairs, where we show that PLMMSE estimation performs only slightly worse than state-of-the-art algorithms, while running an order of magnitude faster. In the dynamic setting, we provide a recursive implementation of the estimator and demonstrate its utility in the context of tracking maneuvering targets from position and acceleration measurements.

Index Terms—Bayesian estimation, minimum mean squared error, linear estimation.

I. INTRODUCTION

Bayesian estimation is concerned with the prediction of a random quantity X based on a set of observations Y , which are statistically related to X . It is well known that the estimator minimizing the mean squared error (MSE) is given by the conditional expectation $\hat{X} = \mathbb{E}[X|Y]$. There are various scenarios, however, in which the minimal MSE (MMSE) estimator cannot be used. This can either be due to implementation constraints, because of the fact that no closed form expression for $\mathbb{E}[X|Y]$ exists, or due to lack of complete knowledge of the joint distribution of X and Y . In these cases, one often resorts to linear estimation. The appeal of the linear MMSE (LMMSE) estimator is rooted in the fact that it possesses an easily implementable closed form expression, which merely requires knowledge of the joint first- and second-order moments of X and Y .

This work was supported in part by the Israel Science Foundation under Grant no. 1081/07, by the European Commission in the framework of the FP7 Network of Excellence in Wireless COMMunications NEWCOM++ (contract no. 216715) and by a Google Research Award.

T. Michaeli and Y. C. Eldar are with the department of electrical engineering and D. Sigalov is with the interdepartmental program for applied mathematics, Technion-Israel Institute of Technology, Haifa, Israel 32000 (phone: +972-4-8294798, +972-4-8294682, fax: +972-4-8295757, e-mail: tomermic@tx.technion.ac.il, dansigal@tx.technion.ac.il, yonina@ee.technion.ac.il). Y. C. Eldar is also a Visiting Professor at Stanford, CA USA.

For example, the amount of computation required for calculating the MMSE estimate of a jump-Markov Gaussian random process from its noisy version grows exponentially in time [1]. By contrast, the LMMSE estimator in this setting possesses a simple recursive implementation, similar to the Kalman filter [2]. A similar problem arises in the area of sparse representations, in which the use of sparsity-inducing Gaussian mixture priors and of Laplacian priors is very common. The complexity of calculating the MMSE estimator under the former prior is exponential in the vector's dimension, calling for approximate solutions [3]. The MMSE estimator under the latter prior does not possess a closed form expression [4], which has motivated the use of alternative estimation strategies such as the maximum a-posteriori (MAP) method.

In practical situations, the reasons for not using the MMSE estimator may only apply to a subset of the measurements. In these cases, it may be desirable to construct an estimator that is linear in part of the measurements and nonlinear in the rest. Partially linear estimation was studied in the statistical literature in the context of regression [5]. In this line of research, it is assumed that the conditional expectation $g(y, z) = \mathbb{E}[X|Y = y, Z = z]$ is linear in y . The goal, then, is to approximate $g(x, y)$ from a set of examples $\{x_i, y_i, z_i\}$ drawn independently from the joint distribution of X , Y and Z . In this paper, our goal is to derive the separable partially linear MMSE (PLMMSE) estimator. Namely, we do not make any assumptions on the structure of the MMSE estimate $\mathbb{E}[X|Y, Z]$, but rather look for the estimator that minimizes the MSE among all functions $g(x, y)$ of the form $Ay + b(z)$.

We show that in certain sparse approximation scenarios, the PLMMSE solution may be computed much more efficiently than the MMSE estimator. We demonstrate the usefulness of the sparse PLMMSE both in static and in dynamic estimation settings. In the static case, we apply our method to the problem of image deblurring from blurred/noisy image pairs [6]. Here, we show that PLMMSE estimation performs only slightly worse than state-of-the-art methods, but is much faster. In the dynamic regime, we provide a recursive implementation of the PLMMSE solution and demonstrate its usefulness in tracking a maneuvering target from position and acceleration measurements. We show the advantage of PLMMSE filtering over state-of-the-art algorithms when the measurements are prone to faults or contain outliers.

The paper is organized as follows. In Section II we present the PLMMSE estimator and discuss some of its properties. In Section III, we show that the PLMMSE method is optimal in a minimax sense among all estimators that solely rely on the second-order statistics of X and Y . In Section IV we derive

the PLMMSE estimator for recovering a signal, sparse in a unitary dictionary, from a pair of observations, one blurred and one noisy. In Section V we apply our method to the problem of image enhancement from blurred/noisy measurement pairs. In Section VI we apply PLMMSE estimation to tracking maneuvering targets.

II. PARTIALLY LINEAR ESTIMATION

We denote random variables (RVs) by capital letters. The pseudo-inverse of a matrix \mathbf{A} is denoted by \mathbf{A}^\dagger . The mean $\mathbb{E}[X]$ of an RV X is denoted μ_X and the auto-covariance matrix $\text{Cov}(X) = \mathbb{E}[(X - \mu_X)(X - \mu_X)^T]$ of X is denoted Γ_{XX} . Similarly, Γ_{XY} stands for the cross-covariance matrix $\text{Cov}(X, Y) = \mathbb{E}[(X - \mu_X)(Y - \mu_Y)^T]$ of two RVs X and Y . The joint cumulative distribution function of X and Y is written as $F_{XY}(x, y) = \mathbb{P}(X \leq x, Y \leq y)$, where the inequalities are element-wise. By definition, the marginal distribution of X is $F_X(x) = F_{XY}(x, \infty)$. In our setting, X is the quantity to be estimated and Y and Z are two sets of measurements thereof. The RVs X , Y and Z take values in \mathbb{R}^M , \mathbb{R}^N and \mathbb{R}^Q , respectively. The MSE of an estimator \hat{X} of X is defined as $\mathbb{E}[\|\hat{X} - X\|^2]$.

We begin by considering the most general form of a partially linear estimator of X based on Y and Z , which is given by

$$\hat{X} = \mathbf{A}(Z)Y + b(Z). \quad (1)$$

Here $\mathbf{A}(z)$ is a matrix-valued function and $b(z)$ is a vector-valued function, so that the realization z of Z is used to choose one of a family of linear estimators of x based on y .

Theorem 1 *Consider estimators of X having the form (1), for some (Borel measurable) functions $\mathbf{A} : \mathbb{R}^Q \rightarrow \mathbb{R}^{M \times N}$ and $b : \mathbb{R}^Q \rightarrow \mathbb{R}^M$. Then the estimator minimizing the MSE within this class is given by*

$$\hat{X} = \mathbf{\Gamma}_{XY|Z}\mathbf{\Gamma}_{YY|Z}^\dagger(Y - \mathbb{E}[Y|Z]) + \mathbb{E}[X|Z], \quad (2)$$

where $\mathbf{\Gamma}_{XY|Z} = \mathbb{E}[(X - \mathbb{E}[X|Z])(Y - \mathbb{E}[Y|Z])^T|Z]$ denotes the cross-covariance of X and Y given Z and $\mathbf{\Gamma}_{YY|Z} = \mathbb{E}[(Y - \mathbb{E}[Y|Z])(Y - \mathbb{E}[Y|Z])^T|Z]$ is the auto-covariance of Y given Z . \blacksquare

Proof: See Appendix A. \blacksquare

Note that (2) is indeed of the form of (1) with $\mathbf{A}(Z) = \mathbf{\Gamma}_{XY|Z}\mathbf{\Gamma}_{YY|Z}^\dagger$ and $b(Z) = \mathbb{E}[X|Z] - \mathbf{\Gamma}_{XY|Z}\mathbf{\Gamma}_{YY|Z}^\dagger\mathbb{E}[Y|Z]$. As can be seen, although the MMSE solution among the class of estimators (1) has a simple form, it requires knowing the conditional covariance $\mathbf{\Gamma}_{XY|Z}$, which limits its applicability. In particular, this solution cannot be applied in cases where we merely know the unconditional covariance Γ_{XY} .

To relax this restriction, we next consider *separable* partially linear estimation. Namely, we seek to minimize the MSE among all functions of the form

$$\hat{X} = \mathbf{A}Y + b(Z), \quad (3)$$

where \mathbf{A} is a deterministic matrix and $b(z)$ is a vector-valued function.

Theorem 2 *Consider estimators of X having the form (3), for some matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ and (Borel measurable) function*

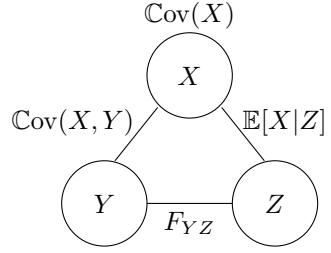


Fig. 1: The statistical knowledge required for computing the PLMMSE estimator (4).

$b : \mathbb{R}^Q \rightarrow \mathbb{R}^M$. Then the estimator minimizing the MSE within this class is given by

$$\hat{X}^{\text{PL}} = \mathbf{\Gamma}_{XY}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger\tilde{Y} + \mathbb{E}[X|Z], \quad (4)$$

where

$$\tilde{Y} = Y - \mathbb{E}[Y|Z]. \quad (5)$$

Proof: See Appendix B. \blacksquare

Note again that (4) is of the form of (3) with $\mathbf{A} = \mathbf{\Gamma}_{XY}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger$ and $b(Z) = \mathbb{E}[X|Z] - \mathbf{\Gamma}_{XY}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger\mathbb{E}[Y|Z]$. The major advantage of this solution with respect to the non-separable estimator (1), is that the only required knowledge regarding the statistical relation between X and Y is of second-order type. Specifically, as we show in Appendix C, (4) can be equivalently written as

$$\begin{aligned} \hat{X}^{\text{PL}} = & \left(\mathbf{\Gamma}_{XY} - \mathbf{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{Y}_Z^{\text{NL}}} \right) \left(\mathbf{\Gamma}_{YY} - \mathbf{\Gamma}_{\hat{Y}_Z^{\text{NL}}\hat{Y}_Z^{\text{NL}}} \right)^\dagger \left(Y - \hat{Y}_Z^{\text{NL}} \right) \\ & + \hat{X}_Z^{\text{NL}}, \end{aligned} \quad (6)$$

where we denoted $\hat{X}_Z^{\text{NL}} = \mathbb{E}[X|Z]$ and $\hat{Y}_Z^{\text{NL}} = \mathbb{E}[Y|Z]$. Therefore, all we need to know in order to be able to compute the separable PLMMSE estimator (4) is the covariance matrix $\mathbf{\Gamma}_{XY}$, the conditional expectation $\mathbb{E}[X|Z]$ and the marginal joint cumulative distribution function F_{YZ} of Y and Z . This is illustrated in Fig. 1. In fact, as we show in Section III, in addition to being optimal among all partially linear methods, the PLMMSE solution (4) is also optimal in a minimax sense among all estimation strategies that rely solely on the quantities appearing in Fig. 1.

The intuition behind (4) is similar to that arising in dynamic estimation schemes, such as the Kalman filter. Specifically, we begin by constructing the estimate $\mathbb{E}[X|Z]$ of X based on the measurements Z , which minimizes the MSE among all functions of Z . Next, we would like to account for Y . However, since Z has already been accounted for, we first need to subtract from Y all variations caused by Z . This is done by constructing the RV \tilde{Y} of (5), which can be thought of as the *innovation* associated with the measurements Y with respect to the initial estimate $\mathbb{E}[X|Z]$. Finally, since we want an estimate that is partially linear in Y , we update our initial estimate with the LMMSE estimate of X based on \tilde{Y} .

Before discussing the minimax-optimality of the PLMMSE estimator, it is insightful to examine several special cases, as we do next.

a) *Independent measurements*: Consider first the case in which Y and Z are statistically independent. In this setting, $\tilde{Y} = Y - \mu_Y$ and therefore the PLMMSE estimator (4) becomes

$$\hat{X}^{\text{PL}} = \boldsymbol{\Gamma}_{XY}\boldsymbol{\Gamma}_{YY}^\dagger(Y - \mu_Y) + \mathbb{E}[X|Z] = \hat{X}_Y^{\text{L}} + \hat{X}_Z^{\text{NL}} - \mu_X, \quad (7)$$

where \hat{X}_Y^{L} denotes the LMMSE estimate of X from Y . Thus, in this setting, the PLMMSE estimate reduces to a linear combination of the LMMSE estimate \hat{X}_Y^{L} and the MMSE estimate \hat{X}_Z^{NL} . The need for subtracting the mean of X arises because both \hat{X}_Y^{L} and \hat{X}_Z^{NL} account for it. Indeed, note that $\mathbb{E}[\hat{X}_Y^{\text{L}}] = \mathbb{E}[\hat{X}_Z^{\text{NL}}] = \mu_X$, so that without subtraction of μ_X , the estimate \hat{X}^{PL} would be biased, with a mean of $2\mu_X$.

b) *Z is independent of X and Y*: Suppose next that both X and Y are statistically independent of Z . Thus, in addition to the fact that $\tilde{Y} = Y - \mu_Y$, we also have $\mathbb{E}[X|Z] = \mu_X$. Consequently, the PLMMSE solution (4) reduces to the LMMSE estimate of X given Y :

$$\hat{X} = \boldsymbol{\Gamma}_{XY}\boldsymbol{\Gamma}_{YY}^\dagger(Y - \mu_Y) + \mu_X = \hat{X}_Y^{\text{L}}. \quad (8)$$

c) *Y is uncorrelated with X and independent of Z*:

Consider the situation in which Y and Z are statistically independent and X and Y are uncorrelated. Then $\tilde{Y} = Y - \mu_Y$, and also $\boldsymbol{\Gamma}_{X\tilde{Y}} = \boldsymbol{\Gamma}_{XY} = 0$ so that (4) becomes the MMSE estimate of X from Z :

$$\hat{X} = \mathbb{E}[X|Z] = \hat{X}_Z^{\text{NL}}. \quad (9)$$

d) *X is independent of Z*: In situations where X and Z are statistically independent, one may be tempted to conclude that the PLMMSE estimator should not be a function of Z . However, this is not necessarily the case. Specifically, although the second term in (4) becomes the constant $\mathbb{E}[X|Z] = \mu_X$ in this setting, it is easily verified that $\boldsymbol{\Gamma}_{X\tilde{Y}} = \boldsymbol{\Gamma}_{XY} = 0$, so that the first term in (4) does not vanish unless X is uncorrelated with Y . As a consequence, the PLMMSE estimator can be written as

$$\hat{X} = \boldsymbol{\Gamma}_{XY}\boldsymbol{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger Y + \mu_X - \boldsymbol{\Gamma}_{XY}\boldsymbol{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \mathbb{E}[Y|Z], \quad (10)$$

in which the last term is a function of Z . This should come as no surprise, though, because if, for instance, $Y = X + Z$, then the optimal estimate is $\hat{X} = Y - Z$, even if X and Z are independent. This solution is clearly a function of Z .

e) *X is uncorrelated with Y*: A similar phenomenon occurs when X and Y are uncorrelated. Indeed in this case, $\boldsymbol{\Gamma}_{X\tilde{Y}} = -\boldsymbol{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{Y}_Z^{\text{NL}}}$, so that the first term in (4) does not vanish unless \hat{X}_Z^{NL} is uncorrelated with \hat{Y}_Z^{NL} . Consequently, the estimator (4) can be expressed as

$$\hat{X} = -\boldsymbol{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{Y}_Z^{\text{NL}}}\boldsymbol{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger Y + \boldsymbol{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{Y}_Z^{\text{NL}}}\boldsymbol{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \mathbb{E}[Y|Z] + \mathbb{E}[X|Z], \quad (11)$$

in which the first term is clearly a linear function of Y .

f) *Additive noise*: Perhaps the most widely studied measurement model corresponds to linear distortion and additive noise. Specifically, suppose that

$$Y = \mathbf{H}X + U, \quad Z = \mathbf{G}X + V, \quad (12)$$

where $\mathbf{H} \in \mathbb{R}^{N \times M}$ and $\mathbf{G} \in \mathbb{R}^{Q \times M}$ are given matrices and U and V are zero-mean RVs such that X , U and V are mutually independent. As we show in Section IV, there are situations in which the distribution of X is such that the complexity of computing the MMSE estimator $\mathbb{E}[X|Y, Z]$ is huge, whereas the complexity of computing $\mathbb{E}[X|Z]$ is modest. In these cases one may prefer to resort to PLMMSE estimation. This method does not correspond to a convex combination of the LMMSE estimate of X from Y and the MMSE estimate of X from Z , as might be suspected. Indeed, substituting $Y = \mathbf{H}X + U$, we have that $\boldsymbol{\Gamma}_{XY} = \boldsymbol{\Gamma}_{XX}\mathbf{H}^T$ and $\boldsymbol{\Gamma}_{YY} = \mathbf{H}\boldsymbol{\Gamma}_{XX}\mathbf{H}^T + \boldsymbol{\Gamma}_{UU}$. Furthermore, $\mathbb{E}[Y|Z] = \mathbf{H}\mathbb{E}[X|Z]$, so that $\boldsymbol{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{Y}_Z^{\text{NL}}} = \boldsymbol{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{X}_Z^{\text{NL}}}\mathbf{H}^T$ and $\boldsymbol{\Gamma}_{\hat{Y}_Z^{\text{NL}}\hat{Y}_Z^{\text{NL}}} = \mathbf{H}\boldsymbol{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{X}_Z^{\text{NL}}}\mathbf{H}^T$. Consequently, the PLMMSE estimator (6) becomes

$$\hat{X} = \mathbf{A}Y + (\mathbf{I} - \mathbf{A}\mathbf{H})\mathbb{E}[X|Z], \quad (13)$$

where \mathbf{I} is the identity matrix and \mathbf{A} is given by

$$\begin{aligned} \mathbf{A} = & \left(\boldsymbol{\Gamma}_{XX} - \boldsymbol{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{X}_Z^{\text{NL}}} \right) \mathbf{H}^T \\ & \times \left(\mathbf{H} \left(\boldsymbol{\Gamma}_{XX} - \boldsymbol{\Gamma}_{\hat{X}_Z^{\text{NL}}\hat{X}_Z^{\text{NL}}} \right) \mathbf{H}^T + \boldsymbol{\Gamma}_{UU} \right)^\dagger. \end{aligned} \quad (14)$$

We see that, as opposed to a convex combination of \hat{X}_Z^{NL} and \hat{X}_Y^{L} , the PLMMSE method reduces to a combination of \hat{X}_Z^{NL} and Y . Furthermore, the weights of this combination are matrices rather than scalars.

As a toy example demonstrating this, suppose that X is a scalar binary RV taking the values ± 1 with equal probability, that $H = G = 1$, and that $U \sim \mathcal{N}(0, \sigma_U^2)$ and $V \sim \mathcal{N}(0, \sigma_V^2)$. It is easily verified that in this case

$$\hat{X}_Z^{\text{NL}} = \mathbb{E}[X|Z] = \frac{\mathcal{N}(Z-1; 0, \sigma_V^2) - \mathcal{N}(Z+1; 0, \sigma_V^2)}{\mathcal{N}(Z-1; 0, \sigma_V^2) + \mathcal{N}(Z+1; 0, \sigma_V^2)}, \quad (15)$$

where $\mathcal{N}(\gamma; \mu, \sigma^2)$ denotes the Gaussian density function with mean μ and variance σ^2 , evaluated at γ . Similarly,

$$\hat{X}_Y^{\text{L}} = \frac{\sigma_{XY}}{\sigma_Y^2}(Y - \mu_Y) + \mu_X = \frac{1}{1 + \sigma_U^2}Y, \quad (16)$$

where we used the facts that $\sigma_Y^2 = \sigma_X^2 + \sigma_U^2$ and $\sigma_{XY} = \sigma_X^2 = 1$. The PLMMSE estimator (13), is therefore given by

$$\hat{X}^{\text{PL}} = \gamma Y + (1 - \gamma)\hat{X}_Z^{\text{NL}}, \quad (17)$$

where $\gamma = (1 - \sigma_{\hat{X}_Z^{\text{NL}}}^2)/(1 + \sigma_U^2 - \sigma_{\hat{X}_Z^{\text{NL}}}^2)$ (see (14)). Figure 3 compares the MSE attained by the PLMMSE method to that of the naive convex-combination estimator

$$\hat{X}^{\text{naive}} = \alpha\hat{X}_Y^{\text{L}} + (1 - \alpha)\hat{X}_Z^{\text{NL}}, \quad (18)$$

for all $\alpha \in [0, 1]$. As can be seen, when $\sigma_U = \sigma_V$, the MMSE of the PLMMSE method is roughly 12% lower than the lowest MSE of the naive estimator. This advantage becomes less significant as σ_U and σ_V draw apart. As mentioned above, though, in multi-dimensional problems the PLMMSE method uses matrix weights rather than scalars, so that its potential for improvement over the naive estimator is yet greater.

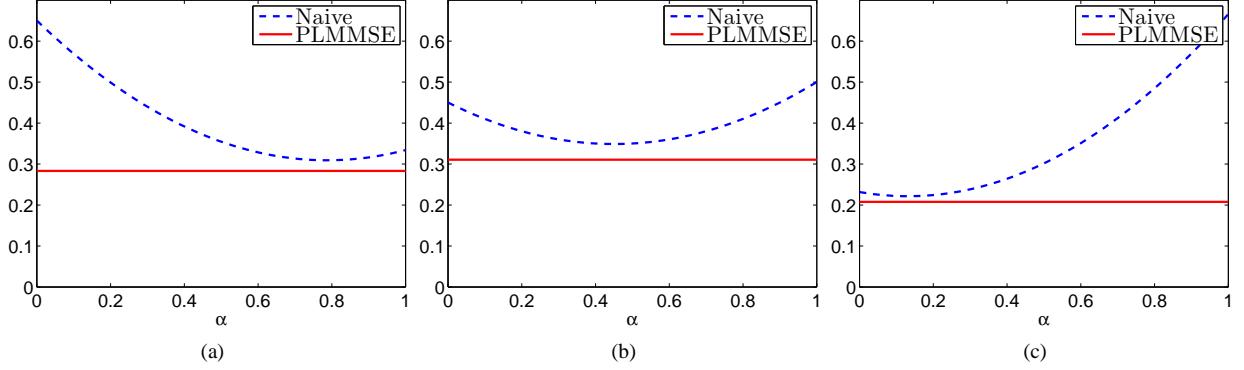


Fig. 2: The MSE attained by \hat{X}^{PL} of (17) and by \hat{X}^{naive} of (18) as a function of α for several values of σ_U and σ_V . (a) $\sigma_U^2 = 0.5$, $\sigma_V^2 = 2$. (b) $\sigma_U^2 = 1$, $\sigma_V^2 = 1$. (c) $\sigma_U^2 = 2$, $\sigma_V^2 = 0.5$.

III. PARTIAL KNOWLEDGE OF STATISTICAL RELATIONS

As discussed in Section II, one of the appealing properties of the PLMMSE solution is that it does not require knowing the entire joint distribution of X and Y , but rather only its second-order moments. However, the fact that the PLMMSE estimator is merely determined by $\mathbb{E}[X|Z]$, $\text{Cov}(X, Y)$ and $F_{YZ}(y, z)$, does not yet imply that it is optimal among all methods that rely solely on these quantities. The question of optimality of an estimator with respect to partial knowledge regarding the joint distribution of the signal and measurements was recently addressed in [7]. One of the notions of optimality considered there, which we adopt here as well, follows from a worst-case perspective. Specifically, any estimator $\hat{X} = g(Y, Z)$, may attain high MSE under certain distributions $F_{XYZ}(x, y, z)$ consistent with our knowledge and it may attain low MSE under other such distributions. We consider an estimator minimax-optimal if its worst-case MSE over the set of all feasible distributions is minimal. For example, it was shown in [7] that the LMMSE estimator \hat{X}_Y^L attains the minimal possible worst-case MSE over the set of distributions $F_{XY}(x, y)$ with given first- and second-order moments.

In the next theorem we show that the PLMMSE method is optimal in the sense that its worst-case MSE over the set of all distributions $F_{XYZ}(x, y, z)$ complying with the knowledge appearing in Fig. 1 is minimal.

Theorem 3 *Let \mathcal{A} be the set of probability distributions of (X, Y, Z) satisfying*

$$\begin{aligned} \text{Cov}(X) &= \Gamma_{XX}, \quad \text{Cov}(X, Y) = \Gamma_{XY}, \quad \mathbb{E}[X|Z] = g(Z), \\ F_{XYZ}(\infty, y, z) &= F_{YZ}(y, z), \end{aligned} \quad (19)$$

where Γ_{XX} and Γ_{XY} are given matrices, $g(z)$ is a given function and $F_{YZ}(y, z)$ is a given cumulative distribution function. Then, among all estimators of X based on Y and Z , the PLMMSE method (4) has the minimal worst-case MSE

$$\sup_{F_{XYZ} \in \mathcal{A}} \mathbb{E}_{F_{XYZ}} \left[\left\| X - \hat{X} \right\|^2 \right], \quad (20)$$

over the set \mathcal{A} .

Proof: See Appendix D. ■

IV. PLMMSE ESTIMATION OF SPARSE VECTORS

We now demonstrate the usefulness of the PLMMSE estimator in the context of sparse approximations. Specifically, consider the situation in which X is known to be sparsely representable in a unitary dictionary $\Psi \in \mathbb{R}^{M \times M}$ in the sense that

$$X = \Psi W \quad (21)$$

for some RV W that is sparse with high probability. More concretely, we assume that the elements of W are given by

$$W_i = S_i B_i, \quad i = 1, \dots, M, \quad (22)$$

where the RVs $\{B_i\}$ and $\{S_i\}$ are statistically independent, $B_i \sim \mathcal{N}(0, 1)$ and $S_i = 0$ (or take small values) with high probability. This assumption is very common in the sparse approximation literature. For example, in [8] and [9] the variables S_i are assumed to follow a Gamma distribution. Here, we assume, as in [3], that they are binary, such that $\mathbb{P}(S_i = \sigma_{1,i}) = 1 - \mathbb{P}(S_i = \sigma_{2,i}) = p_i$ with some $\sigma_{1,i} \geq 0$, $\sigma_{2,i} \geq 0$ and $p_i \geq 0$. In particular, setting $\sigma_{2,i} = 0$ and p_i small corresponds to vectors W that are sparse with high probability.

Assume X is observed through two linear systems, as in (12), where \mathbf{H} is an arbitrary matrix, \mathbf{G} is an orthogonal matrix satisfying $\mathbf{G}^T \mathbf{G} = \alpha^2 \mathbf{I}$ for some constant $\alpha \neq 0$, and U and V are Gaussian RVs with $\Gamma_{UU} = \sigma_U^2 \mathbf{I}$ and $\Gamma_{VV} = \sigma_V^2 \mathbf{I}$. A practical image enhancement scenario and a target tracking situation corresponding to this setting are detailed in sections V and VI, respectively. This setting can be cast in the standard sparse approximation form as

$$\begin{pmatrix} Y \\ Z \end{pmatrix} = \begin{pmatrix} \mathbf{H} \\ \mathbf{G} \end{pmatrix} X + \begin{pmatrix} U \\ V \end{pmatrix}. \quad (23)$$

It is well known that the expression for the MMSE estimate $\mathbb{E}[X|Y, Z]$ in this case generally comprises 2^M summands, which correspond to all different possibilities of sparsity patterns in W [3]. This renders computation of the MMSE estimate prohibitively expensive even for modest values of M and consequently various approaches have been devised to approximate this solution by a small number of terms (see e.g., [3] and references therein). For example, the fast Bayesian

matching pursuit (FBMP) algorithm developed in [3] employs a search in the tree representing all sparsity patterns in order to choose the terms participating in the approximation. We note that FBMP, as well as other sparse recovery methods, can operate with general measurement and dictionary matrices.

There are some special cases, however, in which the MMSE estimate possesses a simple structure, which can be implemented efficiently. As we show next, one such case is when both the channel's response and the dictionary over which X is sparse correspond to orthogonal matrices. As in our setting Ψ is unitary and \mathbf{G} is orthogonal, this implies that we can efficiently compute the MMSE estimate $\mathbb{E}[X|Z]$ of X from Z . Therefore, instead of resorting to schemes for approximating $\mathbb{E}[X|Y, Z]$, we can employ the PLMMSE estimator of X based on Y and Z , which possesses a closed form expression (see (13)) in this situation. This technique is particularly effective when the SNR of the observation Y is much worse than that of Z , since the MMSE estimate $\mathbb{E}[X|Y, Z]$ in this case is close to being partially linear in Y . Such a setting is demonstrated in Section IV-C.

A. MMSE Estimate of a Sparse Signal in a Unitary Dictionary

In our setting

$$Z = \mathbf{G}X + V = \mathbf{G}\Psi W + V, \quad (24)$$

with W of (22). Since \mathbf{G} and Ψ are orthogonal, they are invertible, so that

$$\tilde{Z} = \frac{1}{\alpha} \Psi^T \mathbf{G}^T Z \quad (25)$$

carries the same information on X as Z does, namely

$$\mathbb{E}[X|Z] = \mathbb{E}[X|\tilde{Z}] = \Psi \mathbb{E}[W|\tilde{Z}]. \quad (26)$$

Now, for every $i = 1, \dots, M$, we have that $\tilde{Z}_i = \alpha W_i + \tilde{V}_i$, where $\tilde{V} = \alpha^{-1} \Psi^T \mathbf{G}^T V$ is distributed $\mathcal{N}(0, \sigma_V^2 \mathbf{I})$. Therefore, the set $\{\tilde{Z}_j\}_{j \neq i}$ is statistically independent of the pair (W_i, \tilde{Z}_i) and consequently

$$\begin{aligned} \mathbb{E}[W_i|\tilde{Z}] &= \mathbb{E}[W_i|\tilde{Z}_i] \\ &= \mathbb{E}[W_i|\tilde{Z}_i, S_i = \sigma_{1,i}] \mathbb{P}(S_i = \sigma_{1,i}|\tilde{Z}_i) \\ &\quad + \mathbb{E}[W_i|\tilde{Z}_i, S_i = \sigma_{2,i}] \mathbb{P}(S_i = \sigma_{2,i}|\tilde{Z}_i). \end{aligned} \quad (27)$$

Under the event $S_i = \sigma_{j,i}$ with a fixed $j \in \{1, 2\}$, the RVs W_i and \tilde{Z}_i are jointly normally distributed with mean zero, implying that

$$\mathbb{E}[W_i|\tilde{Z}_i, S_i = \sigma_{j,i}] = \frac{\text{Cov}(W_i, \tilde{Z}_i)}{\text{Cov}(\tilde{Z}_i)} = \frac{\alpha \sigma_{j,i}^2}{\alpha^2 \sigma_{j,i}^2 + \sigma_V^2} \tilde{Z}_i. \quad (28)$$

Finally, using Bayes rule, the term $\mathbb{P}(S_i = \sigma_{1,i}|\tilde{Z}_i)$ reduces to

$$\begin{aligned} &\frac{f_{\tilde{Z}_i|S_i}(\tilde{Z}_i|S_i = \sigma_{1,i})p_i}{f_{\tilde{Z}_i|S_i}(\tilde{Z}_i|S_i = \sigma_{1,i})p_i + f_{\tilde{Z}_i|S_i}(\tilde{Z}_i|S_i = \sigma_{2,i})(1-p_i)} \\ &= \frac{\mathcal{N}(\tilde{Z}_i; 0, \alpha^2 \sigma_{1,i}^2 + \sigma_V^2)p_i}{\mathcal{N}(\tilde{Z}_i; 0, \alpha^2 \sigma_{1,i}^2 + \sigma_V^2)p_i + \mathcal{N}(\tilde{Z}_i; 0, \sigma_{2,i}^2 + \sigma_V^2)(1-p_i)} \end{aligned} \quad (29)$$

and, similarly, $\mathbb{P}(S_i = \sigma_{2,i}|\tilde{Z}_i)$ is given by

$$\begin{aligned} &\mathcal{N}(\tilde{Z}_i; 0, \alpha^2 \sigma_{2,i}^2 + \sigma_V^2)(1-p_i) \\ &\quad \mathcal{N}(\tilde{Z}_i; 0, \alpha^2 \sigma_{1,i}^2 + \sigma_V^2)p_i + \mathcal{N}(\tilde{Z}_i; 0, \sigma_{2,i}^2 + \sigma_V^2)(1-p_i). \end{aligned} \quad (30)$$

Substituting (30) and (28) into (27) leads to the following observation.

Theorem 4 *The MMSE estimate of X of (21) given Z of (24) is*

$$\mathbb{E}[X|Z] = \Psi \tilde{f} \left(\frac{1}{\alpha} \Psi^T \mathbf{G}^T Z \right), \quad (31)$$

where $\tilde{f}(\tilde{z}) = (f(\tilde{z}_1), \dots, f(\tilde{z}_M))^T$, with

$$\begin{aligned} &f(\tilde{z}_i) = \\ &\tilde{z}_i \left(\frac{\alpha \sigma_{1,i}^2 p_i \mathcal{N}(\tilde{z}_i; 0, \alpha^2 \sigma_{1,i}^2 + \sigma_V^2)}{\alpha^2 \sigma_{1,i}^2 + \sigma_V^2} + \frac{(1-p_i) \alpha \sigma_{2,i}^2 \mathcal{N}(\tilde{z}_i; 0, \alpha^2 \sigma_{2,i}^2 + \sigma_V^2)}{\alpha^2 \sigma_{2,i}^2 + \sigma_V^2} \right) \\ &\quad p_i \mathcal{N}(\tilde{z}_i; 0, \alpha^2 \sigma_{1,i}^2 + \sigma_V^2) + (1-p_i) \mathcal{N}(\tilde{z}_i; 0, \alpha^2 \sigma_{2,i}^2 + \sigma_V^2). \end{aligned} \quad (32)$$

Therefore, if, for example, Ψ is a wavelet basis and $\mathbf{G} = \mathbf{I}$ (so that $\alpha = 1$), then $\mathbb{E}[X|Z]$ can be efficiently computed by taking the wavelet transform of Z (multiplication by Ψ^T), applying a scalar shrinkage function on each of the coefficients (namely calculating $f(\tilde{z}_i)$ for the i th coefficient) and applying the inverse wavelet transform (multiplication by Ψ) on the result. Note that the shrinkage curve (32) is different than the soft-threshold operation, originally proposed in [10]. The latter can be obtained as the MAP solution with a Laplacian prior, whereas our function corresponds to the MMSE solution with a Gaussian mixture prior.

B. PLMMSE Estimate of a Sparse Signal

Equipped with a closed form expression for $\mathbb{E}[X|Z]$, we can now obtain an expression for the PLMMSE estimator (13). Specifically, we have that

$$\Gamma_{XX} = \Psi \Gamma_{WW} \Psi^T, \quad (33)$$

where Γ_{WW} is a diagonal matrix with $(\Gamma_{WW})_{i,i} = p_i \sigma_{1,i}^2 + (1-p_i) \sigma_{2,i}^2$. Similarly,

$$\Gamma_{\tilde{X}_Z^{\text{NL}} \tilde{X}_Z^{\text{NL}}} = \Psi \text{Cov}(\tilde{f}(\tilde{Z})) \Psi^T, \quad (34)$$

where $\text{Cov}(\tilde{f}(\tilde{Z}))$ is a diagonal matrix whose (i, i) element is $\beta_i = \text{Cov}(f(\tilde{Z}_i))$. This is due to the fact that the elements of \tilde{Z} are statistically independent and the fact that the function $\tilde{f}(\cdot)$ operates element-wise on its argument. Therefore, the PLMMSE estimator is given in our setting by equation (13) with $\mathbb{E}[X|Z]$ of (31) and with the matrix

$$\begin{aligned} \mathbf{A} &= \Psi \left(\Gamma_{WW} - \text{Cov}(\tilde{f}(\tilde{Z})) \right) \Psi^T \mathbf{H}^T \\ &\quad \times \left(\mathbf{H} \Psi \left(\Gamma_{WW} - \text{Cov}(\tilde{f}(\tilde{Z})) \right) \Psi^T \mathbf{H}^T + \sigma_U^2 \mathbf{I} \right)^\dagger. \end{aligned} \quad (35)$$

Observe that there is generally no closed form expression for the scalars $\beta_i = \text{Cov}(f(\tilde{Z}_i))$, rendering it necessary to compute them numerically. Since each β_i is the variance of a scalar RV, it can be computed very efficiently, either by

approximating the corresponding integral by a sum over a set of points on the real line or by Monte Carlo simulation. In Section V we demonstrate how this can be done in a practical scenario.

An important special case corresponds to the setting in which $p_i = p$, $\sigma_{1,i}^2 = \sigma_1^2$, and $\sigma_{2,i}^2 = \sigma_2^2$ for every i . In this situation, we also have that $\beta_i = \beta$ for every i . Furthermore,

$$\mathbf{\Gamma}_{WW} = (p\sigma_1^2 + (1-p)\sigma_2^2)\mathbf{I} \quad (36)$$

so that \mathbf{A} is simplified to

$$\mathbf{A} = \mathbf{H}^T \left(\mathbf{H}\mathbf{H}^T + \frac{\sigma_U^2}{p\sigma_1^2 + (1-p)\sigma_2^2 - \beta} \mathbf{I} \right)^\dagger. \quad (37)$$

As can be seen, in this setting \mathbf{A} does not involve multiplication by Ψ or Ψ^T . Thus, if \mathbf{H} corresponds to a convolution operation, then \mathbf{A} also corresponds to a filter, which can be efficiently applied in the Fourier domain.

C. Numerical Study

We now compare via simulations the MSE attained by \hat{X}^{PL} to that attained by \hat{X}_Z^{NL} , \hat{X}_Y^{L} and the approximation to $\mathbb{E}[X|Y, Z]$ produced by the FBMP method. Since we generate the signal X and measurements Y and Z according to the assumed model, we do not compare our method to other Bayesian approaches, such as Bayesian compressive sensing (BSC) [8] and sparse Bayesian learning (SBL) [9], which assume a different generative model. Nevertheless, we note that a practical scenario, which deviates from the assumptions of all these methods, was studied in [3] and showed that the performance of FBMP is commonly better than that of BSC and SBL. In terms of running time, FBMP is typically an order of magnitude faster than SBL and roughly twice as slow as BSC.

In our experiment $\Psi \in \mathbb{R}^{256 \times 256}$ was taken to be a Hadamard matrix with normalized columns. The matrix \mathbf{H} corresponded to (circular) convolution with the sequence $h[n] = \exp\{-|n|/8.5\}$ and \mathbf{G} was taken to be diagonal. To comply with the assumption made in [3] that the columns of the measurement matrix are normalized, we normalized the columns of \mathbf{H} to be of norm 0.99 and set $\mathbf{G} = 0.01\mathbf{I}$. We set $p_i = p$, $\sigma_{1,i}^2 = \sigma_1^2$, and $\sigma_{2,i}^2 = 0$ for every i , so that X was truly sparse with high probability. Figure 3 depicts the MSE of all estimators as a function of the input SNR, which we define as $10 \log_{10}(p\sigma_1^2/\sigma^2)$. As can be seen, the MSE of the PLMMSE method is significantly lower than that of \hat{X}_Z^{NL} and \hat{X}_Y^{L} and is very close to that attained by the FBMP method. At low SNR levels and low sparsity levels (high p) the performance of the PLMMSE method is even slightly better than the FBMP approach.

The average running time of the PLMMSE method was 0.6msec for all tested values of p . The average running times of the FBMP method were 52.7msec, 79.6msec and 125.2msec, respectively, for $p = 1/3$, $p = 1/2$ and $p = 2/3$. The ratio between the computational cost of the two approaches, which was two orders of magnitude in this experiment, becomes higher as the dimension of X is increased. At certain dimensions, such as images of size 512×512 (in which

case $M = 512^2$), the FBMP method becomes impractical to apply while PLMMSE estimation can still be used very efficiently.

A word of caution is in place, though. In situations in which the SNR of the measurement Y is roughly the same as that of Z (or better), the FBMP method is advantageous in terms of performance. Therefore in this regime, decision on the use of PLMMSE estimation boils down a performance-complexity tradeoff.

V. APPLICATION TO IMAGE DEBLURRING WITH BLURRED/NOISY IMAGE PAIRS

When taking photos in dim light using a hand-held camera, there is a tradeoff between noise and motion blur, which can be controlled by tuning the shutter speed. Indeed, when using a long exposure time, the image typically comes out blurred due to camera shake. On the other hand, with a short exposure time (and high camera gain), the image is very noisy. In [6] it was demonstrated how a high quality image can be constructed by properly processing two images of the same scene, one blurred and one noisy.

We now show how the PLMMSE approach can be applied in this setting to obtain plausible recoveries at a speed several orders of magnitude faster than any other sparsity-based method. In our setting X , Y and Z correspond, respectively, to the original, blurred (and slightly noisy) and noisy images. Thus, the measurement model is that described by (23), where \mathbf{H} corresponds to spatial convolution with some blur kernel, $\mathbf{G} = \mathbf{I}$, and U and V correspond to white Gaussian images with small and large variances respectively. We further assume that the image X is sparse in some orthogonal wavelet basis Ψ , such that it can be written as in (21) and (22).

As we have seen, in this setting, the PLMMSE estimator can be computed in two stages. In the first stage, we calculate $\hat{X}_Z^{\text{NL}} = \mathbb{E}[X|Z]$ (namely, denoise the image Z) by computing the wavelet transform $\tilde{Z} = \Psi^T Z$, applying the scalar shrinkage function (32) on each wavelet coefficient, and taking the inverse wavelet transform of the result. This stage requires knowledge of the parameters $\{p_i\}$, $\{\sigma_{1,i}\}$, $\{\sigma_{2,i}\}$ and σ_V . To this end, we assume that $\sigma_{2,i} = 0$ for all i (a truly sparse image) and that p_i and $\sigma_{1,i}$ are the same for wavelets coefficients at the same level. In other words, all wavelet coefficients of the noisy image Z at level ℓ correspond to independent draws from the Gaussian mixture

$$f_{\tilde{Z}_i}(\tilde{z}) = p^\ell \mathcal{N}(\tilde{z}; 0, \alpha^2 \sigma_{1,\ell}^2 + \sigma_V^2) + (1-p^\ell) \mathcal{N}(\tilde{z}; 0, \sigma_V^2). \quad (38)$$

Consequently, p^ℓ , $\sigma_{1,\ell}$ and σ_V can be estimated by expectation maximization (EM). In our experiments, we assumed that σ_V is known and thus did not estimate it.

In the second stage, the denoised image \hat{X}_Z^{NL} needs to be combined with the blurred image Y using (13) with \mathbf{A} of (35). As discussed in Section IV-B, this can be carried out very efficiently if $p_i = p$ and $\sigma_{1,i} = \sigma_1$ for all i . For the sake of efficiency¹, we therefore abandon the assumption that

¹The exact solution involving (35) can be computed by using iterative techniques for matrix inversion, in which each iteration comprises filtering operations and forward and inverse wavelet transforms. However, we found that in most cases this approach leads to improvement of only 0.2dB-0.6dB in PSNR and is much more demanding computationally.

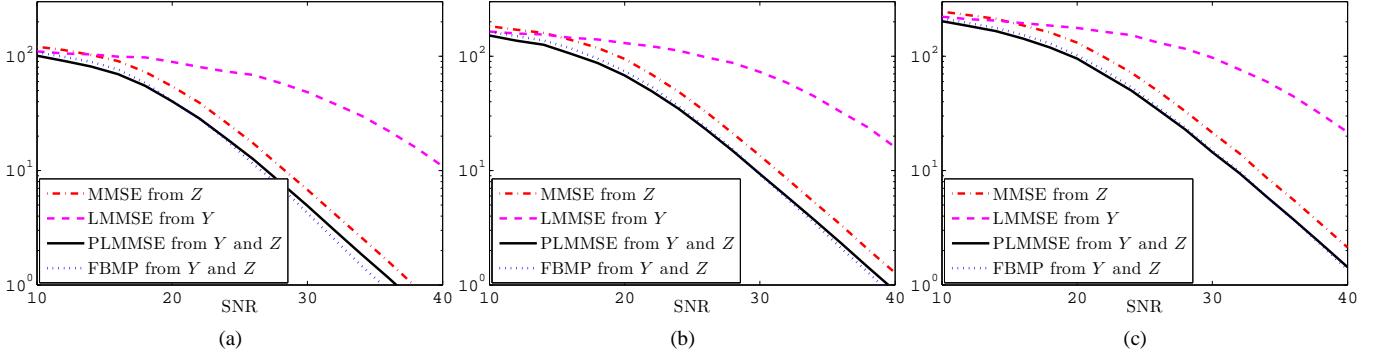


Fig. 3: The MSE attained by \hat{X}_Z^{NL} , \hat{X}_Y^{L} , \hat{X}^{PL} and the approximation of $\mathbb{E}[X|Y, Z]$ produced by the FBMP method [3]. (a) $p = 1/3$. (b) $p = 1/2$. (c) $p = 2/3$.

p_i and $\sigma_{1,i}$ vary across wavelet levels and assume henceforth that all wavelet coefficients are independent and identically distributed. In this case, \mathbf{A} corresponds to the filter

$$A(\omega) = \frac{(\sigma_W^2 - \beta)H^*(\omega)}{(\sigma_W^2 - \beta)|H(\omega)|^2 + \sigma_U^2}, \quad (39)$$

where $H(\omega)$ is the frequency response of the blur kernel. Consequently, the final PLMMSE estimate corresponds to the inverse Fourier transform of

$$\hat{X}_{\text{PLMMSE}}^{\text{F}}(\omega) = \frac{(\sigma_W^2 - \beta)H^*(\omega)Y^{\text{F}}(\omega) + \sigma_U^2\hat{X}_Z^{\text{F}}(\omega)}{(\sigma_W^2 - \beta)|H(\omega)|^2 + \sigma_U^2}, \quad (40)$$

where $Y^{\text{F}}(\omega)$ and $\hat{X}_Z^{\text{F}}(\omega)$ denote the Fourier transforms of Y and \hat{X}_Z^{NL} , respectively. In our experiment, we assumed that the blur $H(\omega)$ and noise variance σ_U^2 are known. In practice, they can be estimated from Y and Z , as proposed in [6]. This stage also requires knowing the scalars $\sigma_W^2 = \mathbb{E}[W^2]$ and $\beta = \mathbb{E}[f^2(\tilde{z})]$, which we estimate as

$$\widehat{\sigma_W^2} = \frac{1}{M} \sum_{i=1}^M \tilde{z}_i^2 - \sigma_V^2, \quad \widehat{\beta} = \frac{1}{M} \sum_{i=1}^M f^2(\tilde{z}_i). \quad (41)$$

Figure 4 demonstrates our approach on the 512×512 Gold-hill image. In this experiment, the blur corresponded to a Gaussian kernel with standard deviation 3.2. To model a situation in which the noise in Y is due only to quantization errors, we chose $\sigma_U = 1/\sqrt{12} \approx 0.3$ and $\sigma_V = 45$. These parameters correspond to a peak signal to noise ratio (PSNR) of 25.08dB for the blurred image and 15.07dB for the noisy image.

We used the orthogonal Symlet wavelet of order 4 and employed 10 EM iterations to estimate p^ℓ and $\sigma_{1,\ell}^2$ in each wavelet level. The entire process takes 1.1 seconds on a Dual-Core 3GHz computer with un-optimized Matlab code. We note that our approach can be viewed as a smart combination of Wiener filtering for image deblurring and wavelet thresholding for image denoising, which are among the simplest and fastest methods available. Consequently, the running time is at least an order of magnitude faster than any other sparsity-based method, including the Bayesian approaches FBMP [3], BCS [8] and SBL [9] and fast ℓ_1 minimization algorithms such as NESTA [11], GPSR [12] and Bergman iterations [13]. As an

example, the authors of [3] reported that FBMP requires a runtime of 38 minutes to recover a 128×128 image from a few thousands of measurements and that GPSR requires 2.7 minutes for the same task. BCS [8] was reported to require 15 seconds for reconstructing a 512×512 image from a few thousands of samples.

As can be seen in Figure 4, the quality of the recoveries corresponding to the denoised image \hat{X}_Z^{NL} and deblurred image \hat{X}_Y^{L} is rather poor with respect to the state-of-the-art BM3D denoising method [14] and BM3D deblurring algorithm [15]. Nevertheless, the quality of the joint estimate \hat{X}_{PLMMSE} surpasses each of these techniques alone. The residual deconvolution (RD) algorithm² proposed in [6] for joint deblurring and denoising slightly outperforms the PLMMSE method in terms of recovery error.

A quantitative comparison on several test images is provided in Table I. This comparison shows that the PSNR attained by the PLMMSE method is, on average, 0.3dB higher than BM3D deblurring, 0.4dB higher than BM3D denoising, and 0.8dB lower than RD. In terms of running times, however, our method is, on average, 11 times faster than BM3D deblurring, 16 times faster than BM3D denoising and 18 times faster than RD. Note that RD requires initialization with a denoised version of Z , for which purpose we used the BM3D algorithm. Consequently, the running time reported in the last column of Table I includes the running time of the BM3D denoising method.

VI. APPLICATION TO MANEUVERING TARGET TRACKING

Next, we demonstrate PLMMSE estimation in the context of maneuvering target tracking. Applications in which there is a need to track the kinematic features of a target are ubiquitous. Often, multiple types of measurements are available. In non-cooperative scenarios, these may include range, bearing, elevation, range rate (Doppler), and more [16]. In navigation applications, measurements may include the signals of global navigation satellite systems and inertial sensors (accelerometers and rate gyros). Such sensors are used in satellites as well as in modern cellular phones, tablet computers and

²In our setting this method does not produce ringing effects and thus the additional de-ringing stage proposed in [6] was not applied.



Fig. 4: Deblurring with a blurred/noisy image pair using PLMMSE estimation and RD [6]. (a) Blurred image Y (top left) and noisy image Z (bottom-right). (b) LMMSE-deblurred image \hat{X}_Y^L (top-left) and MMSE-denoised image \hat{X}_Z^{NL} (bottom-right). (c) BM3D-deblurred image (top left) and BM3D-denoised image (bottom-right). (d) Original image X . (e) PLMMSE estimate \hat{X}^{PL} from Y and Z . (f) RD recovery from Y and Z .

vehicles. Measurements of this type can be fused to aid, *e.g.*, autonomous navigation [17] or traffic monitoring [18].

To model the tracking problem one usually defines a state vector $X(k)$ comprising the target kinematic data, which evolves via the following stochastic linear equation

$$X(k+1) = \mathbf{F}_k X(k) + \mathbf{B}_k W(k). \quad (42)$$

Here, $\{W(k)\}$ is a zero-mean white noise sequence satisfying $\text{Cov}(W(k)) = \sigma_W^2 \mathbf{I}$ for all k and $\{\mathbf{F}_k\}$ and $\{\mathbf{B}_k\}$ are known deterministic matrices. For the simplicity of the exposition, we assume that $X(0) = 0$ (modification to other initializations is trivial). Suppose that two sets of measurements of the state are observed, so that

$$\begin{pmatrix} Y(k) \\ Z(k) \end{pmatrix} = \begin{pmatrix} \mathbf{H}_k \\ \mathbf{G}_k \end{pmatrix} X(k) + \begin{pmatrix} U(k) \\ V(k) \end{pmatrix}, \quad (43)$$

where $\{U(k)\}$ and $\{V(k)\}$ are mutually independent zero-mean white noise sequences satisfying $\text{Cov}(U(k)) = \sigma_U^2 \mathbf{I}$ and $\text{Cov}(V(k)) = \sigma_V^2 \mathbf{I}$, and $\{\mathbf{H}_k\}$ and $\{\mathbf{G}_k\}$ are given matrices.

At the n th time instant, the goal is to obtain an estimate $\hat{X}(n)$ of $X(n)$ based on the measurements $\{Y(k)\}_{k=1}^n$ and $\{Z(k)\}_{k=1}^n$. Ideally, we would like our estimation scheme to possess a recursive structure such that $\hat{X}(n)$ is computed from the previous estimate $\hat{X}(n-1)$ and the current measurements $Y(n)$ and $Z(n)$ without needing to store the entire measurement history.

A simple, yet popular method for modeling maneuvering targets is the dynamic multiple-model method [19] in which $W(k)$ follows a Gaussian mixture distribution. In this case, low intensity noise represents the nominal, non-maneuvering motion regime of the target, and high intensity process noise represents abrupt maneuvers characterized by increased model uncertainty, and caused by, *e.g.*, faults in the actuators of an autonomous aerospace system. Unfortunately, the MMSE solution does not admit a recursive implementation [20] in this setting.

One alternative is to resort in these cases to LMMSE estimation, whose recursive implementation is given by the

TABLE I: Performance of deblurring/denoising on several images. Numbers on the left and right of the slash indicate, respectively, PSNR in dB and running time in seconds.

	\hat{X}_Z^{NL}	\hat{X}_Y^{L}	BM3D Denoising	BM3D Deblurring	PLMMSE	RD
Boat (512 × 512)	25.39 / 0.83	23.45 / 0.06	27.85 / 13.52	28.40 / 10.23	28.05 / 0.88	29.22 / 15.31
Lena (512 × 512)	26.93 / 0.73	24.59 / 0.03	29.47 / 13.22	30.58 / 8.90	30.58 / 0.81	31.37 / 15.19
Mandrill (512 × 512)	21.40 / 0.64	20.59 / 0.06	22.72 / 13.58	21.78 / 9.57	22.58 / 0.72	23.30 / 15.58
Peppers (512 × 512)	26.74 / 0.81	24.89 / 0.08	29.49 / 13.14	29.74 / 8.91	29.80 / 0.88	31.52 / 15.03
Mountain (640 × 480)	19.23 / 0.95	17.69 / 0.09	20.11 / 15.24	18.45 / 11.12	20.03 / 1.05	20.42 / 17.47
Frog (621 × 498)	23.23 / 0.94	22.35 / 0.16	24.00 / 16.07	24.40 / 13.37	24.69 / 1.09	24.69 / 21.14
Gold-hill (512 × 512)	25.90 / 0.69	24.26 / 0.06	27.52 / 13.41	28.70 / 9.54	28.82 / 1.09	29.09 / 21.14
Average	24.12 / 0.81	22.55 / 0.08	25.88 / 14.03	26.01 / 10.23	26.31 / 0.89	27.09 / 16.19

Kalman filter. Another option is to employ approximations of the MMSE estimate, which can be computed recursively, such as the interacting-multiple-model (IMM) filter [1]. The performance of these methods tends to depend heavily on the assumption that the measurement noises are Gaussian. When their actual distribution is unknown, their performance deteriorates.

Sometimes, nonetheless, the MMSE estimate can be calculated in an online manner. As shown below, this happens, *e.g.*, when the state evolves according to the *white acceleration model* [21] and available are acceleration measurements. When supplied with two sets of measurements, only one of which allowing recursive MMSE estimation, it may be advantageous to use PLMMSE estimation rather than approximate MMSE solutions. In this case, under some mild conditions, the PLMMSE estimate can be updated recursively, similarly to the Kalman filter.

Suppose that the distribution of $\{V(k)\}$ and $\{W(k)\}$ is such that, for any $k < n$, $\mathbb{E}[W(k)|Z(1), \dots, Z(n)] = \mathbb{E}[W(k)|Z(k+1)]$ and that the RVs $\{\mathbb{E}[W(k)|Z(k+1)]\}$ are uncorrelated. As we discuss in the sequel, this implies that the MMSE estimate $\hat{X}_Z^{\text{NL}}(n)$ can be computed recursively from $\hat{X}_Z^{\text{NL}}(n-1)$ and $Z(n)$. Our goal is to compute the PLMMSE estimate $\hat{X}^{\text{PL}}(n)$ of $X(n)$ from $\{Y(k)\}_{k=1}^n$ and $\{Z(k)\}_{k=1}^n$. To obtain a recursive implementation, it is insightful to examine first the batch PLMMSE solution. To this end, we define

$$X = \begin{pmatrix} X(1) \\ \vdots \\ X(n) \end{pmatrix}, \quad Y = \begin{pmatrix} Y(1) \\ \vdots \\ Y(n) \end{pmatrix}, \quad Z = \begin{pmatrix} Z(1) \\ \vdots \\ Z(n) \end{pmatrix},$$

$$U = \begin{pmatrix} U(1) \\ \vdots \\ U(n) \end{pmatrix}, \quad V = \begin{pmatrix} V(1) \\ \vdots \\ V(n) \end{pmatrix}, \quad W = \begin{pmatrix} W(0) \\ \vdots \\ W(n-1) \end{pmatrix}. \quad (44)$$

Therewith, we have from (42) that

$$X = \Psi W, \quad Y = HX + U, \quad Z = GX + V, \quad (45)$$

where

$$\Psi = \begin{pmatrix} B_0 & 0 & \cdots & 0 \\ F_1 B_0 & B_1 & \cdots & 0 \\ F_2 F_1 B_0 & F_2 B_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \prod_{k=1}^{n-1} F_{n-k} B_0 & \prod_{k=1}^{n-2} F_{n-k} B_1 & \cdots & B_{n-1} \end{pmatrix}, \quad (46)$$

$$H = \text{diag} (H_1 \ H_2 \ \cdots \ H_n), \quad (47)$$

and

$$G = \text{diag} (G_1 \ G_2 \ \cdots \ G_n). \quad (48)$$

Therefore, as we have seen in (13), the batch PLMMSE estimate of X from Y and Z is given in this case by

$$\mathbb{E}[X|Z] + A(Y - H\mathbb{E}[X|Z]) \quad (49)$$

with the matrix A of (35).

Noting that $\mathbb{E}[X|Z] = \Psi \mathbb{E}[W|Z]$ and taking into account our assumption that $\mathbb{E}[W(k-1)|Z] = \mathbb{E}[W(k-1)|Z(k)]$ and the block lower-triangular structure of Ψ , we see that the k th element in the vector $\mathbb{E}[X|Z]$ is given by

$$\hat{X}_Z^{\text{NL}}(k) = F_{k-1} \hat{X}_Z^{\text{NL}}(k-1) + B_{k-1} \mathbb{E}[W(k-1)|Z(k)]. \quad (50)$$

We have thus obtained a recursive computation of $\hat{X}_Z^{\text{NL}}(k)$ from $\hat{X}_Z^{\text{NL}}(k-1)$ and $Z(k)$.

Next, it remains to determine whether the operation of A admits a recursive implementation. Before we do so, we recall that the LMMSE estimate of X from Y , which is given by

$$\hat{X}_Y^{\text{L}} = \Gamma_{XX} H^T (H \Gamma_{XX} H^T + \sigma_U^2 I)^\dagger Y \quad (51)$$

in this case, can be implemented recursively via the Kalman filter. In our setting, $\Gamma_{XX} = \sigma_W^2 \Psi \Psi^T$ and, due to the assumption that $\text{Cov}(\mathbb{E}[W|Z]) = \beta I$, we have that $\Gamma_{\hat{X}_Z^{\text{NL}} \hat{X}_Z^{\text{NL}}} = \beta \Psi \Psi^T = (\beta/\sigma_W^2) \Gamma_{XX}$. Therefore, the matrix A of the PLMMSE estimate reduces from (14) to

$$\begin{aligned} & \left(1 - \frac{\beta}{\sigma_W^2}\right) \Gamma_{XX} H^T \left(\left(1 - \frac{\beta}{\sigma_W^2}\right) H \Gamma_{XX} H^T + \sigma_U^2 I \right)^\dagger \\ & = \Gamma_{XX} H^T \left(H \Gamma_{XX} H^T + \frac{\sigma_W^2 \sigma_U^2}{\sigma_W^2 - \beta} I \right)^\dagger. \end{aligned} \quad (52)$$

We see that this matrix is the same as that appearing in (51), except for the noise variance which is multiplied here by $\sigma_W^2/(\sigma_W^2 - \beta)$. This implies that multiplication by A corresponds to a Kalman filter with higher observation noise.

We conclude that the complete recursive PLMMSE implementation comprises the following steps:

- Initialization:* $\hat{X}_Z^{\text{NL}}(0) = 0$, $\tilde{X}(0) = 0$, $P_0 = 0$.
- Recursion:* For $k = 1, 2, \dots$ perform the routine summarized in Alg. 1.

Algorithm 1 One Cycle of the Recursive PLMMSE.

Input: $Y(k), Z(k), \hat{X}_Z^{\text{NL}}(k-1), \tilde{X}(k-1), \mathbf{P}_{k-1}$

- 1: $\hat{W}_Z^{\text{NL}}(k-1) = \mathbb{E}[W(k-1)|Z(k)]$.
- 2: $\hat{X}_Z^{\text{NL}}(k) = \mathbf{F}_{k-1}\hat{X}_Z^{\text{NL}}(k-1) + \mathbf{B}_{k-1}\hat{W}_Z^{\text{NL}}(k-1)$.
- 3: $\hat{Y}_Z^{\text{NL}}(k) = \mathbf{H}_k\hat{X}_Z^{\text{NL}}(k)$.
- 4: $\tilde{Y}(k) = Y(k) - \hat{Y}_Z^{\text{NL}}(k)$.
- 5: Update $\tilde{X}(k-1), \mathbf{P}_{k-1}$ using $\tilde{Y}(k)$ via a Kalman step:

$$\begin{aligned} \mathbf{P}_k^- &= \mathbf{F}_{k-1}\mathbf{P}_{k-1}\mathbf{F}_{k-1} + \sigma_W^2\mathbf{B}_k\mathbf{B}_k^T \\ \mathbf{K}_k &= \mathbf{P}_k^-\mathbf{H}_k^T \left(\mathbf{H}\mathbf{P}_k^-\mathbf{H}_k^T + \frac{\sigma_W^2\sigma_U^2}{\sigma_W^2 - \beta} \mathbf{I} \right)^{\dagger} \\ \mathbf{L}_k &= (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{A}_{k-1} \\ \tilde{X}(k) &= \mathbf{L}_k\tilde{X}(k-1) + \mathbf{K}_k\tilde{Y}(k) \\ \mathbf{P}_k &= \mathbf{P}_k^-\mathbf{K}_k \left(\mathbf{H}_k\mathbf{P}_k^-\mathbf{H}_k^T + \frac{\sigma_W^2\sigma_U^2}{\sigma_W^2 - \beta} \mathbf{I} \right) \mathbf{K}_k^T \end{aligned}$$

6: $\hat{X}^{\text{PL}}(k) = \hat{X}_Z^{\text{NL}}(k) + \tilde{X}(k)$.

Output: $\hat{X}^{\text{PL}}(k), \hat{X}_Z^{\text{NL}}(k), \tilde{X}(k), \mathbf{P}_k$

A. Example: Tracking a Maneuvering Target From Position and Acceleration Observations

A common way of describing target kinematics is via the *nearly-constant-velocity model*, also known as the *white acceleration model* [21]. Focusing on one dimension for simplicity, and denoting by $P(k)$ the position of the target at time kT , the state vector $X(k) = (P(k) \ P(k-1) \ P(k-2))^T$ in this model evolves according to (42) with

$$\mathbf{F}_k = \begin{pmatrix} 2 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}. \quad (53)$$

If the sampling interval T is small, then, to close extent, $\mathbf{B}_k = (1 \ 0 \ 0)^T$ (see *e.g.*, [21]). The driving noise W_k in this model corresponds to the target's acceleration.

In many applications [22], the target's velocity changes gradually most of the time apart for abrupt transients, which occur every once in a while. This behavior can be described by letting $W_k = S_k B_k$, where $B_k \sim \mathcal{N}(0, 1)$ and $\mathbb{P}(S_k = \sigma_1) = 1 - \mathbb{P}(S_k = \sigma_2) = p$. If σ_1 is much larger than σ_2 and p is small then the mean time between consecutive large acceleration events is large. In the terminology of the multiple model approach mentioned above, each of the components in the Gaussian mixture corresponds to a different dynamic model.

Suppose that we observe noisy measurements $Y(k)$ and $Z(k)$ of the position $P(k)$ and acceleration $W(k-1)$, respectively. These measurements relate to the state vector via (43), with $\mathbf{H}_k = (1 \ 0 \ 0)$ and $\mathbf{G}_k = (1 \ -2 \ 1)$. Indeed, it is easily verified that in this setting $Z(k) = W(k-1) + V(k)$. Consequently, for $n > k$, $\mathbb{E}[W(k)|Z(1), \dots, Z(n)] = \mathbb{E}[W(k)|Z(k+1)] = f(Z(k+1))$ with the function $f(\cdot)$ of (32).

Equipped with the matrices \mathbf{F}_k , \mathbf{B}_k , \mathbf{H}_k , and \mathbf{G}_k , we estimate the state vector of (42) using the recursive PLMMSE method described above. Note that although $X(k)$ comprises

only the target's positions at times k , $k-1$, and $k-2$, estimating the velocity and acceleration is straightforward as $\hat{P}(k) - \hat{P}(k-1)$ and $\hat{P}(k) - 2\hat{P}(k-1) + \hat{P}(k-2)$, respectively.

We compare the performance of the recursive PLMMSE method with the one of a standard KF which provides, recursively, the LMMSE optimal estimate of the state using the measurement sets $\{Z(k)\}$ and $\{Y(k)\}$, as well as with that of the IMM filter [1] which is known to be extremely effective in multiple model estimation problems [19]. The main idea underlying the IMM algorithm is to maintain a bank of primitive Kalman filters, each matched to a different model in the given model set. Each filter produces a local estimate with an associated error covariance using its initial estimate and covariance and the current measurement. The key element of the IMM scheme is the interaction block that generates, using all local estimates, individual initial conditions for each of the primitive filters in the bank. In our case, the two models maintained by the IMM method correspond to the two possible realizations of S_k such that one model corresponds to the low process noise variance representing the nominal target regime and the second model for the maneuvering one.

We simulated a random sequence $\{X(k)\}$ for $k = 1, \dots, 1000$ according to (42) initialized at $X(0) = 0$ and driven by a process noise $\{W(k)\}$ having, at each time, a two-modal Gaussian mixture distribution, with $p = 0.05$, $\sigma_1 = 10$, and $\sigma_2 = 1$. The state position is observed via a Gaussian measurement equation with $\sigma_U = 5$ and the covariance of the Gaussian measurement noise of the acceleration, σ_V , is swept from 1 to 15. Averaged over 500 independent Monte Carlo runs, the average squared position, velocity, and acceleration errors are presented, respectively, on the top, middle, and bottom of the left chart of Fig. 5. Outperformed by the IMM, the PLMMSE method scores better than the Kalman filter since it optimally utilizes the acceleration measurements. It is noticeable that at high values of σ_V^2 , Kalman's and PLMMSE's errors coincide indicating that acceleration measurements do not carry valuable information in addition to that carried by $\{Y(k)\}$.

In many practical scenarios the distribution of the position measurement noise is far from Gaussian (see *e.g.*, [23] and references therein). On the right chart of Fig. 5 we present the position, velocity and acceleration errors obtained for a Gaussian mixture distribution of the position noise having the same first- and second-order statistics as before. Such a distribution may model occasional outlier measurements or sensor faults [24]. None of the filters is supplied with this information and only the first- and second-order moments are provided to the algorithms. Utilizing the position measurements in a linear manner, both Kalman and PLMMSE keep the performance unchanged relative to the Gaussian case. In a contradistinction, the IMM algorithm results in an inferior performance. This phenomenon is tightly related to the statement of Theorem 3 claiming that the PLMMSE method is ensured to attain a smaller worst-case MSE in comparison to any other estimation technique provided the appropriate moments are kept constant.

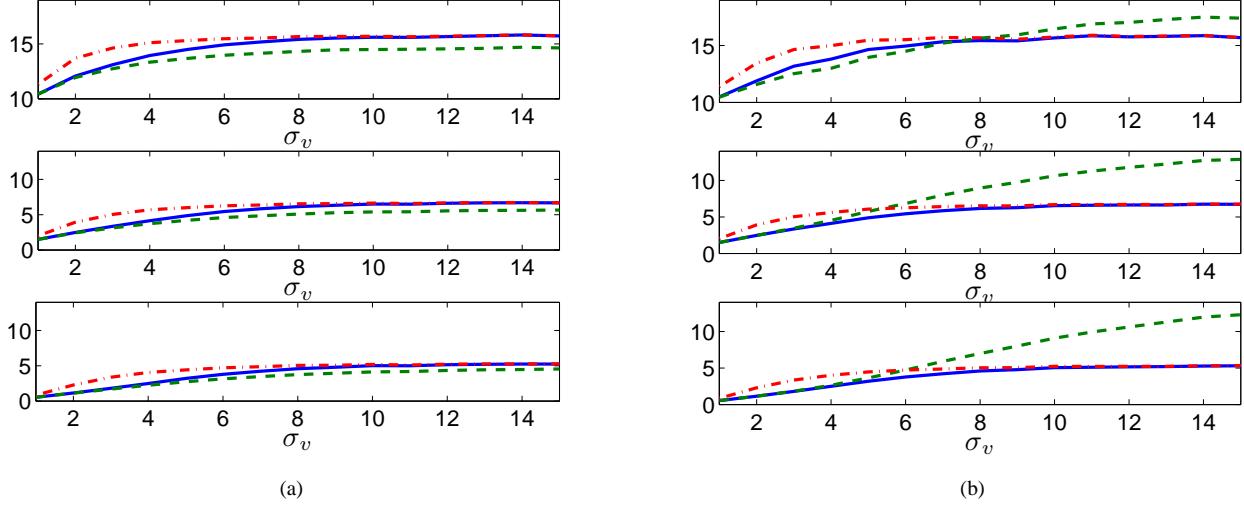


Fig. 5: Mean squared estimation errors of the position (top), velocity (middle) and acceleration (bottom) vs. σ_v using the PLMMSE approach (solid), IMM (dashed) and standard KF (dash-dotted). (a) $\{U(k)\}$ have a Gaussian distribution. (b) $\{U(k)\}$ have a Gaussian mixture distribution.

VII. CONCLUSIONS

In this paper we derived the PLMMSE estimator, which is the method whose MSE is minimal among all functions that are linear in Y . We showed that the PLMMSE solution depends only on the joint second-order statistics of X and Y , which renders it applicable in a wide variety of situations. Furthermore, we showed that this estimator attains the lowest worst-case MSE over the set of distributions whose joint second-order moments of X and Y are fixed. We demonstrated our approach in the context of recovering a vector, which is sparse in a unitary dictionary, from a pair of noisy measurements. In this setting, the PLMMSE solution achieves an MSE very close to that attained by iterative approximation strategies, such as the FBMP method of [3], and is much cheaper computationally. We applied our method to the problems of image enhancement from blurred/noisy image pairs and maneuvering target tracking from position and acceleration measurements. In both applications, we showed that PLMMSE estimation performs close to state-of-the-art algorithms. In the image enhancement setting, we showed that it can run much faster than competing approaches. In the context of target tracking, we demonstrated the insensitivity of the solution to the distribution of the noise in Y . This property provides robustness against sensor faults and outlier measurements, problems which are very common in target tracking situations.

APPENDIX A PROOF OF THEOREM 1

Using the smoothing property, the MSE of any estimator of the form (1) is given by

$$\mathbb{E} \left[\mathbb{E} \left[\|X - \mathbf{A}(Z)Y - b(Z)\|^2 | Z \right] \right]. \quad (54)$$

Thus, for every specific value z that Z can take, the optimal choice of $\mathbf{A}(z)$ and $b(z)$ is that minimizing the inner expectation. The solution to this minimization problem corresponds

to the LMMSE estimate of X based on Y , under the joint distribution of (X, Y) given Z , concluding the proof.

APPENDIX B PROOF OF THEOREM 2

We start by noting that the set \mathcal{B} of RVs constituting candidate estimates is a closed linear subspace within the space of finite-second-order-moment RVs taking values in \mathbb{R}^M . Therefore, the MMSE estimate \hat{X} within this subspace, which is the projection of the RV X onto \mathcal{B} , is the unique³ RV whose estimation error $\hat{X} - X$ is orthogonal to every RV of the form $\mathbf{A}Y + b(Z)$. To demonstrate that \hat{X} of (4) indeed satisfies this property, note that the inner product between $\hat{X} - X$ and $\mathbf{A}Y + b(Z)$ is given by

$$\begin{aligned} \mathbb{E} \left[(\hat{X} - X)^T (\mathbf{A}Y + b(Z)) \right] &= \text{Tr} \left\{ \mathbb{E} \left[(\hat{X} - X)Y^T \right] \mathbf{A}^T \right\} \\ &\quad + \text{Tr} \left\{ \mathbb{E} \left[(\hat{X} - X)b(Z)^T \right] \right\}. \end{aligned} \quad (55)$$

Substituting (4), the expectation within the second term becomes

$$\begin{aligned} \mathbb{E} \left[\left(\Gamma_{X\tilde{Y}} \Gamma_{\tilde{Y}\tilde{Y}}^\dagger \tilde{Y} + \mathbb{E}[X|Z] - X \right) b(Z)^T \right] \\ = \Gamma_{X\tilde{Y}} \Gamma_{\tilde{Y}\tilde{Y}}^\dagger \mathbb{E} \left[\tilde{Y} b(Z)^T \right] + \mathbb{E} \left[(\mathbb{E}[X|Z] - X) b(Z)^T \right]. \end{aligned} \quad (56)$$

Recall that $\tilde{Y} = Y - \mathbb{E}[Y|Z]$ is the estimation error incurred in estimating Y from Z . Consequently, \tilde{Y} and $X - \mathbb{E}[X|Z]$ are uncorrelated with every function of Z and, in particular, with $b(Z)$, so that this expression vanishes. Similarly, substituting (4) and expressing $Y = \tilde{Y} + \mathbb{E}[Y|Z]$, the expectation within

³In an almost-sure sense.

the first summand in (55) becomes

$$\begin{aligned} & \mathbb{E} \left[\left(\mathbf{\Gamma}_{X\tilde{Y}} \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \tilde{Y} + \mathbb{E}[X|Z] - X \right) Y^T \right] \\ &= \mathbf{\Gamma}_{X\tilde{Y}} \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \mathbb{E} \left[\tilde{Y} (\tilde{Y} + \mathbb{E}[Y|Z])^T \right] \\ &\quad - \mathbb{E} \left[(X - \mathbb{E}[X|Z]) (\tilde{Y} + \mathbb{E}[Y|Z])^T \right]. \end{aligned} \quad (57)$$

Being a function of Z , the RV $\mathbb{E}[Y|Z]$ is uncorrelated with \tilde{Y} and $X - \mathbb{E}[X|Z]$ so that this expression can be reduced to

$$\begin{aligned} & \mathbf{\Gamma}_{X\tilde{Y}} \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \mathbb{E} \left[\tilde{Y} \tilde{Y}^T \right] - \mathbb{E} \left[(X - \mathbb{E}[X|Z]) W^T \right] \\ &= \mathbf{\Gamma}_{X\tilde{Y}} \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}} - \mathbb{E} \left[(X - \mu_X + \mu_X - \mathbb{E}[X|Z]) \tilde{Y}^T \right] \\ &= \mathbf{\Gamma}_{X\tilde{Y}} - \mathbf{\Gamma}_{X\tilde{Y}} + \mathbb{E} \left[(\mathbb{E}[X|Z] - \mu_X) \tilde{Y}^T \right] \\ &= \mathbf{\Gamma}_{X\tilde{Y}} - \mathbf{\Gamma}_{X\tilde{Y}} \\ &= 0, \end{aligned} \quad (58)$$

where we used the facts that $\mathbb{E}[\tilde{Y}] = 0$, that $\mathbf{\Gamma}_{X\tilde{Y}} \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}} = \mathbf{\Gamma}_{X\tilde{Y}}$ [25, Lemma 2], and that \tilde{Y} is uncorrelated with $\mathbb{E}[X|Z]$ (due to the same argument as above). This completes the proof.

APPENDIX C DERIVATION OF EQUATION (6)

By definition,

$$\begin{aligned} \mathbf{\Gamma}_{X\tilde{Y}} &= \mathbb{E}[(X - \mu_X)(Y - \mathbb{E}[Y|Z])^T] \\ &= \mathbb{E}[(X - \mu_X)(Y - \mu_Y + \mu_Y - \mathbb{E}[Y|Z])^T] \\ &= \mathbb{E}[(X - \mu_X)(Y - \mu_Y)^T] - \mathbb{E}[(X - \mu_X)(\mathbb{E}[Y|Z] - \mu_Y)^T] \\ &= \mathbf{\Gamma}_{XY} - \mathbb{E}[\mathbb{E}[(X - \mu_X)(\mathbb{E}[Y|Z] - \mu_Y)^T|Z]] \\ &= \mathbf{\Gamma}_{XY} - \mathbb{E}[(\mathbb{E}[X|Z] - \mu_X)(\mathbb{E}[Y|Z] - \mu_Y)^T] \\ &= \mathbf{\Gamma}_{XY} - \mathbf{\Gamma}_{\hat{X}_Z^{\text{NL}} \hat{Y}_Z^{\text{NL}}}, \end{aligned} \quad (59)$$

where $\hat{X}_Z^{\text{NL}} = \mathbb{E}[X|Z]$ and $\hat{Y}_Z^{\text{NL}} = \mathbb{E}[Y|Z]$. Here, the fourth equality is a result of the smoothing property and the last equality follows from the facts that $\mathbb{E}[\mathbb{E}[X|Z]] = \mu_X$ and $\mathbb{E}[\mathbb{E}[Y|Z]] = \mu_Y$. In a similar manner, it is easy to show that

$$\mathbf{\Gamma}_{Y\tilde{Y}} = \mathbf{\Gamma}_{YY} - \mathbf{\Gamma}_{\hat{Y}_Z^{\text{NL}} \hat{Y}_Z^{\text{NL}}}. \quad (60)$$

Using (60) and the fact that \tilde{Y} is uncorrelated with $\mathbb{E}[Y|Z] - \mu_Y$, we obtain

$$\begin{aligned} \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}} &= \mathbb{E}[\tilde{Y} \tilde{Y}^T] \\ &= \mathbb{E}[(Y - \mathbb{E}[Y|Z]) \tilde{Y}^T] \\ &= \mathbb{E}[(Y - \mu_Y) \tilde{Y}^T] - \mathbb{E}[(\mathbb{E}[Y|Z] - \mu_Y) \tilde{Y}^T] \\ &= \mathbf{\Gamma}_{YY} \\ &= \mathbf{\Gamma}_{YY} - \mathbf{\Gamma}_{\hat{Y}_Z^{\text{NL}} \hat{Y}_Z^{\text{NL}}}. \end{aligned} \quad (61)$$

Substituting (59) and (61) into (4) leads to (6).

APPENDIX D PROOF OF THEOREM 3

Let $\varepsilon(F_{XYZ}, \hat{X}) = \mathbb{E}_{F_{XYZ}}[\|\hat{X} - X\|^2]$ denote the MSE incurred by an estimator \hat{X} of X based on Y and Z , when

the joint distribution of X , Y and Z is $F_{XYZ}(x, y, z)$. It is easily verified that

$$\begin{aligned} & \varepsilon(F_{XYZ}, \hat{X}_{\text{PLMMSE}}) = \text{Tr}\{\mathbf{\Gamma}_{XX}\} \\ & - \text{Tr} \left\{ (\mathbf{\Gamma}_{XY} - \mathbf{\Gamma}_{\hat{X}_Z^{\text{NL}} \hat{Y}_Z^{\text{NL}}})(\mathbf{\Gamma}_{YY} - \mathbf{\Gamma}_{\hat{Y}_Z^{\text{NL}} \hat{Y}_Z^{\text{NL}}})^\dagger (\mathbf{\Gamma}_{XY} - \mathbf{\Gamma}_{\hat{X}_Z^{\text{NL}} \hat{Y}_Z^{\text{NL}}})^T \right\} \end{aligned} \quad (62)$$

for all $F_{XYZ} \in \mathcal{A}$. Therefore, in particular, (62) is also the worst-case MSE of \hat{X}^{PL} over \mathcal{A} . We next make use of the following lemma.

Lemma 1 *There exists a distribution F_{XYZ}^* in the set \mathcal{A} of distributions satisfying (19), under which the PLMMSE estimate of X based on Y and Z coincides with the MMSE estimate $\mathbb{E}[X|Y, Z]$.*

Proof: See Appendix E. ■

Now, any estimator \hat{X} that is a function of Y and Z satisfies

$$\begin{aligned} \sup_{F_{XYZ} \in \mathcal{A}} \varepsilon(F_{XYZ}, \hat{X}) &\geq \varepsilon(F_{XYZ}^*, \hat{X}) \\ &\geq \min_{\hat{X}} \varepsilon(F_{XYZ}^*, \hat{X}) \\ &= \varepsilon(F_{XYZ}^*, \mathbb{E}[X|Y, Z]) \\ &= \varepsilon(F_{XYZ}^*, \hat{X}^{\text{PL}}) \\ &= \max_{F_{XYZ} \in \mathcal{A}} \varepsilon(F_{XYZ}, \hat{X}^{\text{PL}}), \end{aligned} \quad (63)$$

where the first line follows from the fact that $F_{XYZ}^* \in \mathcal{A}$, the third line is a result of the fact that the MMSE and PLMMSE estimators coincide under F_{XYZ}^* , and the last line is due to the fact that $\varepsilon(F_{XYZ}, \hat{X}^{\text{PL}})$ is constant as a function of F_{XYZ} over \mathcal{A} . We have thus established that the worst-case MSE of any estimator over \mathcal{A} is greater or equal to the worst-case MSE of the PLMMSE solution over \mathcal{A} , proving that \hat{X}^{PL} is minimax optimal.

APPENDIX E PROOF OF LEMMA 1

We prove the statement by construction. Let Y and Z be two RVs distributed according to F_{YZ} and denote $h(Z) = \mathbb{E}[Y|Z]$ and $\tilde{Y} = Y - h(Z)$. Let U be a zero-mean RV, statistically independent of the pair (\tilde{Y}, Z) , whose covariance matrix is

$$\mathbf{\Gamma}_{UU} = \mathbf{\Gamma}_{XX} - \text{Cov}(g(Z)) - \mathbf{\Gamma}_{X\tilde{Y}} \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \mathbf{\Gamma}_{\tilde{Y}X}. \quad (64)$$

It can be easily verified that this is the covariance matrix of the estimation error of the LMMSE estimate of $\tilde{X} = X - \mathbb{E}[X|Z]$ based on $\tilde{Y} = Y - \mathbb{E}[Y|Z]$. Therefore, this is a valid covariance matrix. Consider the RV⁴

$$X = \mathbf{\Gamma}_{X\tilde{Y}} \mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger \tilde{Y} + g(Z) + U. \quad (65)$$

We will show that the so constructed X , Y and Z satisfy the constraints (19). Indeed, using the fact that U has zero mean and is statistically independent of Z , we find that the conditional expectation of X of (65) given Z is

$$\mathbb{E}[X|Z] = g(Z). \quad (66)$$

⁴Recall that $\mathbf{\Gamma}_{X\tilde{Y}}$ and $\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}$ are functions of $\text{Cov}(X, Y)$ and F_{YZ} , which are given.

Furthermore, since \tilde{Y} , $g(Z)$ and U are pairwise uncorrelated, the covariance of X of (65) can be computed as

$$\begin{aligned}\text{Cov}(X) &= \mathbf{\Gamma}_{X\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}X} + \text{Cov}(g(Z)) + \mathbf{\Gamma}_{UU} \\ &= \mathbf{\Gamma}_{X\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger\mathbf{\Gamma}_{\tilde{Y}X} + \text{Cov}(g(Z)) \\ &\quad + \mathbf{\Gamma}_{XX} - \text{Cov}(g(Z)) - \mathbf{\Gamma}_{X\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger\mathbf{\Gamma}_{\tilde{Y}X} \\ &= \mathbf{\Gamma}_{XX},\end{aligned}\quad (67)$$

where we substituted (64). Finally, the cross covariance of X of (65) and Y is given by

$$\begin{aligned}\text{Cov}(X, Y) &= \mathbf{\Gamma}_{X\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger\mathbf{\Gamma}_{\tilde{Y}Y} + \text{Cov}(g(Z), h(Z)) \\ &= \mathbf{\Gamma}_{X\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}} + \text{Cov}(g(Z), h(Z)) \\ &= \mathbf{\Gamma}_{X\tilde{Y}} + \text{Cov}(g(Z), h(Z)) \\ &= \mathbf{\Gamma}_{XY} - \text{Cov}(g(Z), h(Z)) + \text{Cov}(g(Z), h(Z)) \\ &= \mathbf{\Gamma}_{XY},\end{aligned}\quad (68)$$

where the second equality follows from the third line of (61), the third equality follows from [25, Lemma 2], and the fourth equality follows from (59). Equations (66), (67) and (68) demonstrate that the distribution F_{XYZ}^* associated with X , Y and Z , belongs to the family of distributions \mathcal{A} satisfying (19).

Next, we show that the PLMMSE and MMSE estimators coincide under F_{XYZ}^* . Indeed, since U is statistically independent of the pair (Y, Z) , we have that $\mathbb{E}[U|\tilde{Y}, Z] = \mathbb{E}[U] = 0$, so that

$$\begin{aligned}\mathbb{E}[X|Y, Z] &= \mathbf{\Gamma}_{X\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger\mathbb{E}[\tilde{Y}|Y, Z] + \mathbb{E}[g(Z) + U|Y, Z] \\ &= \mathbf{\Gamma}_{X\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger(Y - h(Z)) + g(Z) + \mathbb{E}[U|\tilde{Y}, Z] \\ &= \mathbf{\Gamma}_{X\tilde{Y}}\mathbf{\Gamma}_{\tilde{Y}\tilde{Y}}^\dagger(Y - h(Z)) + g(Z),\end{aligned}\quad (69)$$

where we used the fact that there is a one-to-one transformation between the pair (Y, Z) and the pair (\tilde{Y}, Z) . This expression is partially linear in Y , implying that this is also the PLMMSE estimator in this setting. Thus, for the distribution F_{XYZ}^* , the PLMMSE estimator is optimal not only among all partially linear functions, but also among *all* functions of Y and Z .

REFERENCES

- [1] H. A. P. Blom and Y. Bar-Shalom, "The interacting multiple model algorithm for systems with Markovian switching coefficients," *IEEE Trans. Autom. Control*, vol. 33, no. 8, pp. 780–783, 1988.
- [2] O. L. V. Costa, "Linear minimum mean square error estimation for discrete-time Markovian jump linear systems," *IEEE Trans. Autom. Control*, vol. 39, no. 8, pp. 1685–1689, 1994.
- [3] P. Schniter, L. C. Potter, and J. Ziniel, "Fast Bayesian matching pursuit," in *Information Theory and Applications Workshop (ITA'08)*, 2008, pp. 326–333.
- [4] M. Girolami, "A variational method for learning sparse and overcomplete representations," *Neural Computation*, vol. 13, no. 11, pp. 2517–2532, 2001.
- [5] W. Härdle and H. Liang, *Statistical Methods for Biostatistics and Related Fields*, 2007, ch. Partially linear models, pp. 87–103.
- [6] L. Yuan, J. Sun, L. Quan, and H. Y. Shum, "Image deblurring with blurred/noisy image pairs," in *ACM SIGGRAPH 2007 papers*. ACM, 2007, pp. 1–10.
- [7] T. Michaeli and Y. C. Eldar, "Hidden relationships: Bayesian estimation with partial knowledge," *IEEE Trans. Signal Process.*, vol. 59, no. 5, 2011.
- [8] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2346–2356, 2008.
- [9] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *The Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.
- [10] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *Journal of the american statistical association*, pp. 1200–1224, 1995.
- [11] S. Becker, J. Bobin, and E. Candès, "NESTA: A fast and accurate first-order method for sparse recovery," *Arxiv preprint arXiv:0904.3367*, 2009.
- [12] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 1, no. 4, pp. 586–597, 2007.
- [13] W. Yin, S. Osher, D. Goldfarb, and J. Darbon, "Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing," *SIAM Journal on Imaging Sciences*, vol. 1, no. 1, pp. 143–168, 2008.
- [14] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [15] ———, "Image restoration by sparse 3d transform-domain collaborative filtering," in *SPIE Electronic Imaging*, 2008.
- [16] X. R. Li and V. P. Jilkov, "A survey of maneuvering target trackingPart III: Measurement models," in *Proc. 2001 SPIE Conf. on Signal and Data Processing of Small Targets*, 2001, pp. 423–446.
- [17] S. Sukkarieh, E. M. Nebot, and H. F. Durrant-Whyte, "A high integrity IMU/GPS navigation loop for autonomous land vehicle applications," *IEEE Trans. Robot. Autom.*, vol. 15, no. 3, pp. 572–578, 1999.
- [18] P. Mohan, V. N. Padmanabhan, and R. Ramjee, "Nericell: rich monitoring of road and traffic conditions using mobile smartphones," in *Proceedings of the 6th ACM conference on Embedded network sensor systems*, 2008, pp. 323–336.
- [19] Y. Bar-Shalom, X. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York : Wiley, 2001.
- [20] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, 2002.
- [21] X. R. Li and V. P. Jilkov, "Survey of maneuvering target tracking. Part I: Dynamic models," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1333–1364, 2003.
- [22] D. Sigalov and Y. Oshman, "Linear optimal state estimation in hybrid dynamical systems with actuator faults," in *Proc. Fusion 2011, 14th International Conference on Information Fusion*, 2011.
- [23] M. Idan and J. L. Speyer, "Cauchy estimation for linear scalar systems," *IEEE Trans. Autom. Control*, vol. 55, no. 6, pp. 1329–1342, 2010.
- [24] E. Mazor, A. Averbuch, Y. Bar-Shalom, and J. Dayan, "Interacting multiple model methods in target tracking: a survey," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 34, no. 1, pp. 103–123, 1998.
- [25] A. Torokhti and P. Howlett, "An optimal filter of the second order," *IEEE Trans. Signal Process.*, vol. 49, no. 5, pp. 1044–1048, 2002.